

IN THE UNITED STATES
PATENT AND TRADEMARK OFFICE

Declaration and Power of Attorney

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name.

I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled **MULTIPROCESSOR SYSTEM WITH CACHE-BASED SOFTWARE BREAKPOINTS** the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by an amendment, if any, specifically referred to in this oath or declaration.

I acknowledge the duty to disclose all information known to me which is material to patentability as defined in Title 37, Code of Federal Regulations, 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

None

I hereby claim the benefit under Title 35, United States Code, 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, 112, I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, 1.56 which became available between the filing date of the prior application and the national or PCT international filing date of this application:

None

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

I hereby appoint the following attorney(s) with full power of substitution and revocation, to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected therewith:

Lester H. Birnbaum	(Reg. No. 25830)
Richard J. Botos	(Reg. No. 32016)
Gerard A. deBlasi	(Reg. No. 34149)
Anthony Grillo	(Reg. No. 36535)
Mark A. Kurisko	(Reg. No. 38944)
Robert P. Marley	(Reg. No. 32914)
Scott W. McLellan	(Reg. No. 30776)
Geraldine Monteleone	(Reg. No. 40097)
Scott J. Rittman	(Reg. No. 39010)
Ferdinand M. Romano	(Reg. No. 32752)
David L. Smith	(Reg. No. 30592)
John P. Veschi	(Reg. No. 39058)

I hereby appoint the attorney(s) on ATTACHMENT A as associate attorney(s) in the aforementioned application, with full power solely to prosecute said application, to make alterations and amendments therein, to receive the patent, and to transact all business in the Patent and Trademark Office connected with the prosecution of said application. No other powers are granted to such associate attorney(s) and such associate attorney(s) are specifically denied any power of substitution or revocation.

Full name of 1st joint inventor: Michael Richard Betker

Inventor's signature Michael Richard Betker Date 2/1/2002

Residence: Allentown, Lehigh County, Pennsylvania

Citizenship: United States of America

Post Office Address: ~~606 Cedar Hills Drive~~
~~Allentown, Pennsylvania 18103~~ MRB 2/1/2002
1627 Harding Circle
Whitehall, PA 18052

Full name of 2nd joint inventor: Han Q. Nguyen

Inventor's signature Han Q. Nguyen Date 2/6/2002

Residence: ~~Marlboro, Burlington County, New Jersey~~
Allentown, Lehigh County, Pennsylvania

Citizenship: United States of America

Post Office Address: ~~2 Tharp Lane~~
~~Marlboro, New Jersey 07746~~ HQN 2/6/2002
815 Dorset Ct.
Allentown, PA 18104

Full name of 3rd joint inventor: Bryan Schlieder

Inventor's signature Bryan Schlieder Date 2/1/2002

Residence: Bethlehem, Northampton County, Pennsylvania

Citizenship: United States of America

Post Office Address: 2935 Hecktown Road
Bethlehem, Pennsylvania 18020

Full name of 4th joint inventor: Shaun Patrick Whalen

Inventor's signature Shaun Patrick Whalen Date 2/1/2002

Residence: Wescosville, Lehigh County, Pennsylvania

Citizenship: United States of America

Post Office Address: 5301 Hanover Drive
Wescosville, Pennsylvania 18106

Full name of 5th joint inventor: Jay Patrick Wilshire

Inventor's signature Jay Patrick Wilshire Date 2/1/2002

Residence: Pennsburg, Bucks County, Pennsylvania

Citizenship: United States of America

Post Office Address: 2030 Miller Road
Pennsburg, Pennsylvania 18073

ATTACHMENT A

Attorney Name(s):	Joseph B. Ryan	Reg. No. 37922
	Kevin M. Mason	Reg. No. 36597
	William E. Lewis	Reg. No. 39274
	Robert J. Mauri	Reg. No. 41180
	Wayne L. Ellenbogen	Reg. No. 43602
	James M. Loeffler	Reg. No. 37873
	James F. Harrington	Reg. No. 44741

Telephone calls should be made to Joseph B. Ryan of Ryan, Mason & Lewis, LLP at:

Phone No.: (516) 759-7517
Fax No.: (516) 759-9512

All written communications are to be addressed to:

Ryan, Mason & Lewis, LLP
90 Forest Avenue
Locust Valley, New York 11560

MULTIPROCESSOR SYSTEM
WITH CACHE-BASED SOFTWARE BREAKPOINTS

Field of the Invention

The present invention relates generally to multiprocessor integrated circuits and other types of systems which include multiple processors, and more particularly to techniques for implementing software breakpoints in such systems.

Background of the Invention

Software breakpoints are typically implemented as assembly instructions embedded in a program. A given one of these assembly instructions when encountered in the program causes execution of the program to transfer to a debugger. As is well known, the use of designated breakpoints in conjunction with the debugger allows a programmer to inspect the state of the program in order to fix a problem or better understand program behavior.

In the case of a single conventional processor, a software breakpoint may be embedded in a program by overwriting an existing instruction at the desired breakpoint location with an assembly instruction in the form of a specified debug operation code, also referred to as an opcode. The overwritten instruction is stored by the debugger so that it may subsequently replace the inserted opcode upon resumption of normal operation, as will be described below. Examples of conventional debug opcodes include opcodes referred to as DEBUG, TRAP, etc. These debug opcodes generally vary depending on the processor.

After a given breakpoint is taken and the programmer is ready to continue with normal execution of the program, a two step process typically takes place. First, the debugger replaces the debug opcode with the previously-overwritten existing instruction, so that this instruction alone may be executed (or "stepped"). Then the debug opcode is restored at the desired breakpoint location so that the next time program execution reaches this point, another breakpoint will occur.

Many high performance processors include an instruction cache. An instruction cache enables the processor to store frequently accessed instructions in a fast local memory that adapts its contents to the executing program.

A multiprocessor system generally includes multiple processors each connected to a common bus. Also connected to the common bus is a shared main memory that stores instructions and data for use by the processors. In such a system, each of the processors will typically have its own instruction cache. These instruction caches are particularly important in multiprocessor systems in that the caches serve to significantly reduce the bus traffic associated with the instruction and data fetches of the multiple processors.

The conventional single-processor software breakpoints described above generally do not work in a multiprocessor system. More particularly, in a multiprocessor system, once a particular one of the processors has taken a software breakpoint, the corresponding actual instruction must be written to the shared memory so the stopped processor can subsequently fetch and execute it. However, while the actual instruction is in shared memory, the other processors may fetch and execute it, thus missing the breakpoint. Thus, in a shared-memory multiprocessor system, the programmer utilizing conventional single-processor software breakpoint techniques would be likely to observe erratic behavior, e.g., sometimes a given processor would stop at a software breakpoint and other times it would not.

Examples of known multiprocessor systems include the DSP16410 and DSP16270 shared-memory multiprocessor digital signal processors (DSPs) available from Agere Systems Inc. of Allentown, Pennsylvania, USA. Each of these systems includes two processors, denoted DSP0 and DSP1. These processors do not have instruction caches, but can execute from a common shared memory. Consider a case where a software breakpoint has been set in the shared memory and executed by DSP0. In order to resume normal execution after the breakpoint is taken, the actual instruction code at the breakpoint location should be placed in local memory that is private to DSP0. This enables DSP0 to execute the actual instruction without preventing DSP1 from hitting the desired breakpoint. Implementing multiprocessor software breakpoints on the DSP16410 or the DSP16270 therefore may require local instruction memory, which adds to the cost of the device. In addition, rebuilding the executable code so that the breakpoint debug opcode is moved to the address space corresponding to the local memory can be inconvenient and time consuming for the programmer. These and other similar problems arise in a wide variety of other conventional multiprocessor systems.

A need therefore exists for improved techniques for implementing software breakpoints in shared-memory multiprocessor systems, so as to avoid one or more of the above-described problems.

Summary of the Invention

5 The invention provides improved techniques for implementing software breakpoints in a shared-memory multiprocessor system.

10 In accordance with one aspect of the invention, a software breakpoint is implemented in a multiprocessor system having a number of processors each coupled to a main memory. Each of the processors preferably has an instruction cache associated therewith. An instruction for which a breakpoint is to be inserted is retrieved from a corresponding instruction address in the main memory, and a breakpoint code, e.g., a debug opcode, is inserted at the instruction address in main memory.

15 The invention may be implemented using a debugger which communicates with the multiprocessor system via an interface. The debugger saves the retrieved instruction that has been replaced with the breakpoint code, and when the breakpoint code is executed by a given one of the processors, a breakpoint is taken and the debugger takes control of the given processor. Before the debugger resumes execution, it stores the retrieved instruction in the instruction cache for the given processor, and sets a use-once indicator associated with the instruction as stored in the corresponding instruction cache for that processor.

20 The use-once indicator, when set for the instruction as stored in the instruction cache, is operative via cache control logic to clear a validity indicator associated with the instruction after a single fetch of the instruction from the instruction cache. As a result, subsequent attempts by the given processor to access the instruction as stored in the instruction cache will cause the processor to retrieve the breakpoint code at the instruction address in main memory.

25 Advantageously, the software breakpoint techniques of the present invention ensure that a specified debug opcode or other breakpoint code can remain present in the main memory of the multiprocessor system at all times, such that all of the processors of the system will reliably fetch and execute that breakpoint code even if one or more of these processors are resuming execution from the breakpoint.

Brief Description of the Drawings

FIG. 1 is a block diagram of an example multiprocessor system in which the present invention is implemented.

FIG. 2 illustrates the operation of an instruction cache of the FIG. 1 multiprocessor system in accordance with the invention.

FIG. 3 illustrates the interaction of the FIG. 1 multiprocessor system with a debugger.

FIG. 4 shows a state and flow diagram illustrating the operation of a multiprocessor system in accordance with the invention.

FIG. 5 is an example illustrating the operation of a three-processor implementation of a multiprocessor system in accordance with the invention.

Detailed Description of the Invention

The invention will be illustrated herein in conjunction with an exemplary shared-memory multiprocessor system. It should be understood, however, that the invention is more generally applicable to any shared-memory multiprocessor system in which it is desirable to provide improved performance through the use of software breakpoints. The term "multiprocessor system" as used herein is intended to include any device in which retrieved instructions are executed using one or more processors. The term "processor" is intended to include, by way of example and without limitation, microprocessors, central processing units (CPUs), very long instruction word (VLIW) processors, single-issue processors, multi-issue processors, digital signal processors (DSPs), application-specific integrated circuits (ASICs), and other types of data processing devices, as well as portions and combinations of these and other devices.

FIG. 1 shows a shared-memory multiprocessor system 100 in which the present invention is implemented. The system 100 includes a number of processors 102-1, 102-2, . . . 102-N, each coupled to a main memory 104 via a common bus 106. Also associated with each of the processors 102- i , $i = 1, 2, \dots N$, is a corresponding multiplexer 108- i and instruction cache (Icache) 110- i . In operation, the processors 102 each retrieve instructions for execution, as well as associated data, from the main memory 104. The instructions may be retrieved directly from the main memory 104 via a first path through the multiplexer 108, or from the instruction cache 110 via a second path

through the multiplexer 108. The processors 102 also return instruction execution results to the main memory 104 via bus 106, although separate return paths from the processors to the bus 106 are not explicitly shown in the figure. These return paths may be combined in whole or in part with the illustrative signal lines used to retrieve the instructions and data, i.e., bidirectional signal lines may be used.

Also included in the multiprocessor system 100 is a set of cache control logic 112. Although not explicitly shown as such in the figure, the cache control logic 112 may be coupled to each of the instruction caches 110. Alternatively, the cache control logic 112 may be distributed across the instruction caches 110 such that a portion of the cache control logic 112 forms an internal portion of each of the instruction caches 110. Other arrangements may also be used, e.g., various combinations of centralized and distributed logic. The operation of the cache control logic 112 will be described in greater detail below in conjunction with the description of FIGS. 2 through 5. Suitable hardware or software elements, or various combinations thereof, may be configured in a straightforward manner to provide the cache control logic functionality described herein, as will be appreciated by those skilled in the art. The cache control logic 112 may therefore be implemented in hardware, in software, or as a combination of hardware and software.

In accordance with the present invention, the instruction caches 110 of the FIG. 1 system are configured so as to permit a debugger to place and mark an existing instruction, for which a breakpoint is to be established, in a given one of the instruction caches of a particular processor. Advantageously, the placing and marking of the instruction ensure that the other processors of the system will not miss the breakpoint when the particular processor is resumed. The illustrative embodiment of the invention implements this capability in part through the addition of a control bit, referred to herein as a use-once (U) bit, to each set of instruction information stored in the instruction cache, and by permitting the debugger to write this information directly to the instruction cache.

FIG. 2 illustrates the instruction processing in a given one of the instruction caches 110 of the FIG. 1 system, i.e., a particular instruction cache 110-*i*, although it should be understood that substantially the same instruction processing is utilized in each of the instruction caches 110. In the illustrative embodiment, when the corresponding processor 102-*i* fetches an instruction from the instruction cache 110-*i*, it presents an address 200 to the cache. The address 200 as shown in FIG.

2 includes a tag field 202, an index field 204, and a block offset field 206. The address 200 in this example is a 32-bit address, with bits 0-3 corresponding to the block offset, bits 4-11 corresponding to the index, and bits 12-31 corresponding to the tag.

It should be emphasized that this particular address arrangement is by way of example only, and other address lengths and field configurations may be used. In other words, the invention does not require that the presented address have a particular format.

The index field 204 of the address 200 is used to select one of several sets in the instruction cache 110-*i*. More particularly, the particular value of the index field 204 specifies a particular set 210 within a group of such sets, where each of the sets corresponds generally to a set of instructions stored in the cache. Each instruction in the set 210 includes a valid (V) bit, the above-mentioned use-once (U) bit, a tag field, and the instruction data.

The instruction cache 110-*i* in accordance with the invention includes logic elements 212 and 214, which may be implemented in a straightforward manner well-known to those skilled in the art. As shown in the figure, logic element 212 compares the tag field 202 of the presented address 200 to the tag fields of the set 210 as selected using the index field 204, and logic element 214 performs an AND operation on the output of the logic element 212 and the valid bit. The output of logic element 214 is at a logic high level if for a given instruction in the set 210 the tag field 202 matches its tag field and its valid bit is set, i.e., at a logic high level. In this situation, the cache is said to "hit" at the instruction address 200, and otherwise to "miss" at the instruction address 200.

In normal operation absent any inserted software breakpoints, if there is a match between the tag field 202 of the address and the tag field of a given one of the instructions, and the valid bit is set for the given instruction, then the desired instruction data is in the cache 110-*i* and is returned to the corresponding processor 102-*i*. If there is no match or the valid bit is clear, the desired instruction data is not in the cache 110-*i*, and therefore must be fetched from the main memory 104.

As indicated above, the use-once bit is utilized to implement software breakpoints in an advantageous manner in accordance with the techniques of the invention. Such breakpoints may be inserted using an arrangement such as that shown in FIG. 3. In this particular software breakpoint processing configuration 300, a debugger 302 interacts via an interface 304 with the multiprocessor 100 of FIG. 1. The debugger 302, which may also be referred to as a debugger tool, may be

implemented using appropriately-configured hardware and software elements, as is well known. The interface 304 may also be implemented in manner well understood by those skilled in the art, and is therefore not described in detail herein. It is to be appreciated that the invention does not require the use of any particular type of debugger or multiprocessor system interface. The debugger will typically run on a processor that is not an element of the system 100, but this is by way of example and not limitation.

In order to program a software breakpoint for a given processor of the multiprocessor system 100, a programmer utilizes the debugger 302 of FIG. 3 to overwrite an existing instruction at the desired breakpoint location in main memory 104 with a specified debug opcode. However, the existing instruction to be overwritten is first fetched from main memory 104 and saved by the debugger for subsequent storage in one or more of the instruction caches 110 upon resumption of execution from a taken breakpoint, as will be described in greater detail below.

When a given processor 102-*i* executes the inserted debug opcode from the main memory 104, the breakpoint is said to be taken by that processor, and execution of the program is transferred to the debugger 302, e.g., via a designated exception, transition to a debug operation mode, etc. After the breakpoint is taken by processor 102-*i* and the programmer is ready to continue with normal execution of the program, the debugger 302 writes the instruction data portion of the appropriate element in the cache set 210 with the corresponding information from the actual instruction for which the breakpoint was taken, and sets the use-once bit for that element in the cache set 210.

The processor 102-*i* resumes execution of the program after the breakpoint is taken by executing a fetch request using the address of the actual instruction for which the breakpoint was taken. This fetch request is serviced from the instruction cache 110-*i*. However, immediately after servicing this request, and in accordance with the techniques of the invention, the set use-once bit causes the cache control logic to clear the valid bit for that element of the set 210. Thus, the next time the processor 102-*i* issues a fetch request for this instruction address, the cache will "miss" and the debug opcode will automatically be fetched again from main memory 104.

Advantageously, the above-described mechanism ensures that the desired debug opcode is always present in main memory 104, such that all of the processors 102 will reliably fetch and execute it even if one or more of these processors are resuming execution from the breakpoint.

It should be noted that it is common to associate attributes with instruction addresses to control the handling of the instruction fetch. For example, one known attribute is the so-called “noncacheable” attribute. This attribute indicates whether a particular instruction address should be allowed to be stored in the instruction cache. Each of the processors 102 typically includes hardware for interpreting such attributes. For an instruction address with the noncacheable attribute, the cache control logic always treats a fetch to this address as a miss.

The above-described use-once bit, however, changes the handling of the noncacheable attribute so that a software breakpoint can be set on a noncacheable instruction. For a software breakpoint on a noncacheable instruction, a given processor 102-*i* fetches the breakpointed instruction, but does not put the instruction in its instruction cache. After the breakpoint is taken by processor 102-*i* and the programmer is ready to continue with normal execution of the program, the debugger 302 writes the instruction data portion of the appropriate element in the cache set 210 with the corresponding information from the actual noncacheable instruction for which the breakpoint was taken, and sets the use-once bit for that element in the cache set 210. The set use-once bit causes the cache control logic to fetch this instruction address from the cache even though its attribute says that it is noncacheable. Once the fetch is serviced from the cache, the valid bit of the instruction is cleared as previously described. Thus, the use-once bit mechanism of the present invention allows software breakpoints to be set on both cacheable and noncacheable instructions, and to be taken by a given processor without impacting other processors in the multiprocessor system 100.

FIG. 4 shows a state and flow diagram 400 which illustrates in greater detail the above-described cache-based software breakpoint process as implemented in multiprocessor system 100. The diagram 400 includes an IDLE state 402, a USE ONCE state 404, and a MISS state 406.

The system is initially in the IDLE state, e.g., upon execution of a reset command as shown. From the IDLE state, a determination is made in step 410 as to whether an instruction fetch has been requested by a given processor 102 of the system. If not, the process remains in the IDLE state. Otherwise, step 412 determines if there has been an instruction cache hit for the given processor.

If there is no instruction cache hit in step 412, the given processor is stalled and the instruction request is directed to the main memory 104 of the system, as indicated in step 414, and the process then enters the MISS state. If there is an instruction cache hit in step 412, step 416 determines if the USE ONCE mode is set, i.e., determines if the use-once bit for the instruction cache is set. If the use-once bit is set, the given processor is stalled and the valid and use-once bits of the instruction cache are cleared, as indicated in step 418, and the process then enters the USE ONCE state. If the use-once bit is not set in step 416, the instruction fetch data is returned from the instruction cache as indicated in step 420.

From the USE ONCE state, the process in step 422 removes the processor stall that was implemented in step 418, and proceeds to step 420 to return the instruction fetch data from the instruction cache.

From the MISS state, the process in step 424 determines if the instruction fetch data has been returned from the main memory. If the instruction fetch data has not been returned, the process remains in the MISS state. If the instruction fetch data has been returned, it is returned to the requesting processor in step 426, and the process returns to the IDLE state.

At least a portion of the operations described in conjunction with FIG. 4 can be implemented using the cache control logic 112 of FIG. 1. For example, the cache control logic 112 may be configured to clear the valid bit in step 418 based on the determination of the state of the use-once bit in step 416. As indicated previously, suitable arrangements of hardware, software or both for accomplishing these and other operations in the cache control logic 112 will be readily apparent to those skilled in the art.

It is also to be appreciated that the particular processing operations shown in FIG. 4 are by way of example only, and should not be construed as limitations of the invention. Those skilled in the art will recognize that the techniques of the invention can be implemented using other arrangements of processing operations.

FIG. 5 is a more particular example illustrating the operation of a three-processor implementation of the multiprocessor system 100 in accordance with the invention. The three processors are denoted P0, P1 and P2 for purposes of this example. The figure includes portions (a) through (g) which represent the changes in the state of the three-processor system 100 as software

breakpoints are taken with the use-once mechanism described above. For each of the portions (a) through (g), the states of the three processors are shown. Each line indicates a different aspect of the state of the corresponding processor. The P_i state indicates whether the corresponding processor is in a debug state or a nondebug state. These states are also referred to as modes. The other states utilized in the figure are as follows:

P_i cache[...] V state indicates the state of the valid bit for a specific set in the cache for P_i .

P_i cache[...] U state indicates the state of the use-once bit for a specific set in the cache for

P_i .

P_i cache[...] Tag state indicates the state of the tag part for a specific set in the cache for P_i .

P_i cache[...] Data state indicates the state of the data part for a specific set in the cache for

P_i .

P_i cache[...] V state indicates the state of the valid bit for a specific set in the cache for P_i .

The expression $\text{index}(x)$ denotes the index field 204 of the address of x . The expression $\text{cache}[y]$ denotes the set of cache 210 addressed by y . The expression $\text{tag}(z)$ indicates the tag field 202 of the address of z .

In (a) all three processors are executing operational code and are not in the debug mode.

In (b), P_0 has fetched an instruction on which a software breakpoint has been set. The instruction on which the software breakpoint is set is indicated by swbpt . This instruction is stored in the cache in the set $\text{index}(\text{swbpt})$. After this instruction has been fetched into the cache, the valid bit V of this set is 1 and the use-once bit U of this set is 0. The tag 202 and data parts of this set are written with the tag part of the address and the actual data fetched from main memory 104. The data for this instruction is the DEBUG opcode.

In (c), P_0 has executed the DEBUG opcode and entered its debug mode. The state of the cache set holding the DEBUG instruction is the same as in (b).

In (d), P_0 is ready to exit its debug mode and resume normal operation. The debugger 302 has written the use-once bit to 1 and the data part of the cache set to the actual instruction that was replaced by a DEBUG opcode in main memory 104. Meanwhile, P_2 has fetched the same

instruction on which the software breakpoint is set. P2 fetches the instruction from main memory 104, which is still the DEBUG opcode, and writes it into its instruction cache, setting the valid bit, tag, and data parts of the corresponding cache set. This state illustrates how the use-once logic allows one processor to execute the actual instruction on which the software breakpoint is set without interfering with another processor hitting the breakpoint.

In (e), P0 has resumed normal operation. When P0 resumed, it addressed the cache with the address of swbpt. Since the V and U bits for this set are 1, the cache returned the bopcode as the fetched instruction and cleared the V and U bits. P2 has executed the DEBUG opcode and entered debug mode.

In (f), P0 continues normal operation. P2 is ready to exit its debug mode and resume normal operation. The debugger 302 has written the use-once bit to 1 and the data part of the cache set to the actual instruction that was replaced by a DEBUG opcode in main memory 104.

In (g), P2 has resumed normal operation. When P2 resumed, it addressed the cache with the address of swbpt. Since the V and U bits for this set are 1, the cache returned the bopcode as the fetched instruction and cleared the V and U bits. P0 is still executing normally.

The foregoing example is intended for purposes of illustration only, and should not be viewed as limiting the scope of the invention in any way.

Advantageously, the invention provides substantially improved software breakpoint capability in a multiprocessor system with little or no impact on processor cycle time, minimal hardware cost, and without requiring program alteration. Importantly, the action of clearing the valid bit in response to the set use-once bit does not take place during the lookup process where the index is used to find a set and compare the tags. Rather, the action of clearing the valid bit takes place after the hit determination is complete. In addition, there is no need for private local memory in each processor since the existing instruction cache memory is used. Furthermore, the executable code of a given user program can be used as is. There is no need to relocate it to private local memory so that the debug opcode can be managed.

It should also be noted that the invention makes the debugger more efficient when resuming from a breakpoint. With the invention, the debugger no longer has to perform the previously-described two-step process of "stepping" the actual instruction and then replacing the debug opcode.

More particularly, the invention eliminates the latter of these two steps, which is a significant performance advantage for the debugger.

Although illustrated using an actual instruction cache, the invention can be implemented at least in part utilizing dedicated hardware which provides the functionality described above but does not otherwise serve as an actual instruction cache in normal operation of a multiprocessor system. In other words, the dedicated hardware is configured substantially as an instruction cache but is used only in implementing breakpoints in accordance with the techniques of the invention. The term "instruction cache" as used herein is intended to include this type of dedicated hardware.

The present invention may be configured to meet the requirements of a variety of different processing applications and environments, using any desired types and arrangements of processor, instruction cache, bus and main memory elements. The above-described embodiments of the invention are therefore intended to be illustrative only. For example, although the illustrative embodiment utilizes a single-bit use-once indicator for breakpointed instructions stored in an instruction cache, multiple-bit indicators can also be used, as can other arrangements of different types of single-bit or multiple-bit indicators. In addition, the particular manner in which the a use-once indicator is updated may be altered. Furthermore, the breakpoint code, instruction, instruction address and cache set configurations may be varied as required to accommodate a given processing application or environment. These and numerous other alternative embodiments within the scope of the following claims will be apparent to those skilled in the art.